SRv6

Record-Speed Standardization and Deployment

Clarence Filsfils
cf@cisco.com

Thank you

- Lead operators
- EcoSystem Partners
- Academic Partners
- IETF Partners
- Cisco SR Team

segment-routing.net

- Detailed recording of this presentation: link
- SRv6 uSID: training
- SRv6 Stateless Slicing: demo
- SRv6 Ultra-Scale SR Policy: 26 uSID push at linerate: demo
- BGP PIC Edge with SRv6 Summarization: ISIS UPA: demo
- Path Tracing: demo

Deployment Status

Record-Speed Deployment

- 3 years of commercial deployment (2019-2022)
- ~100M SRv6 subscribers
- ~100 deployments, with ~14 public reports
- Across markets (Web, SP, Enterprise) and geographies (Asia, EU, US)

SoftBank	Rakuten	Indosat	MTN Uganda	Noia
Iliad	Alibaba	China Telecom	China Bank	Cernet2
Free	Bell Canada	China Unicom	Line	

NEW

Alibaba - Full Stack SRv6 Deployment

- *"Predictable network services to every single application, based on the full stack SRv6 innovations across endpoints(container,host,P4 gateway), network devices and controller/network service center"* YuanChao Su, Alibaba
- SID ⇔ Se
- Massive I
- Seamless
  - Many u

Co-Development

✓ SRv6 Service Anchors on C8000

✓ 2. The uSID POC in C8000

ork'

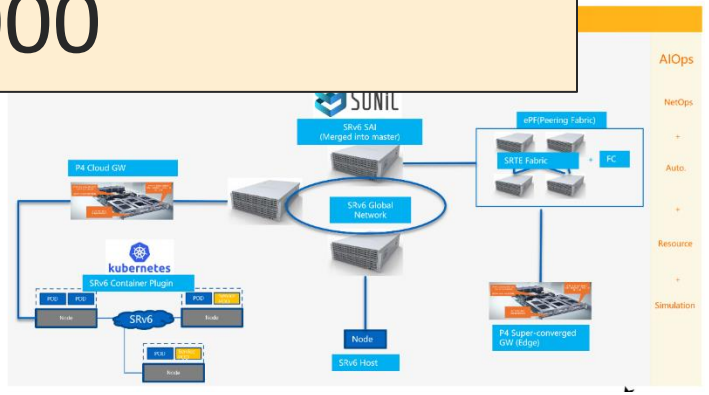

Alibaba: Full stack SRv6 towards a 'Predictable Network'

Commitment to SR Lead-Operators

Standard-Based Technology

Vendor eco-system

Open-Source eco-system

Mature Standardization

Much faster standardization than usual

Sign of the SRv6 Industry Endorsement

- Proposed Standard
  - RFC 8402 SR Architecture
  - RFC 8754 SRv6 DataPlane
  - RFC 8986 SRv6 Network Programming
  - RFC xxxx SRv6 ISIS Extension
  - RFC xxxx SRv6 BGP Extension
  - RFC xxxx SRv6 OAM
  - RFC xxxx SR Policy

Rich Eco-System

- 25 HW implementations
  - Cisco Systems, Nokia, Arrcus, Kaloom, NoviFlow, Huawei, ZTE, Juniper
  - Broadcom, Barefoot, Intel SmartNIC, Marvell, Mellanox,
  - Spirent, Ixia
  - Multiple Interop Reports
- 11 open-source platforms/ Applications
  - Linux, FD.io VPP, P4, iptables, nftables, snort, SERA, ExaBGP, GoBGP, GoBMP, Contiv-VPP

Many Successful Interops

- 2020/04: EANTC: SRv6 interop between Cisco, Huawei, Juniper, Arrcus, Ixia (link)
  - Classic IPv6 nodes as SRv6 transit nodes
  - SRv6-L3VPN for IPv4 and IPv6 services
  - SRv6 TI-LFA FRR link protection with SRH insert
  - SRv6 EVPN for E-Line and EVPN L3VPN services
  - SRv6 TE SR Policy
- 2021/02: NetOne Systems (link)
  - Cisco XR, Cisco NX, Juniper
- 2021/09: EANTC: SRv6 interop between Cisco, Huawei, Juniper, Nokia, Spirent (link)
  - SRv6-Based Global IPv4 and IPv6 services
  - SRv6-L3VPN for IPv4 and IPv6 services
  - SRv6 TI-LFA FRR local SRLG protection with SRH insert
  - SRv6 EVPN for E-Line and EVPN L3VPN services
  - IGP Flex-Algo using TWAMP-measured link delays
- 2022: EANTC: Cisco did not participate due to the COVID situation

Stateless & Scalable Network Partitioning

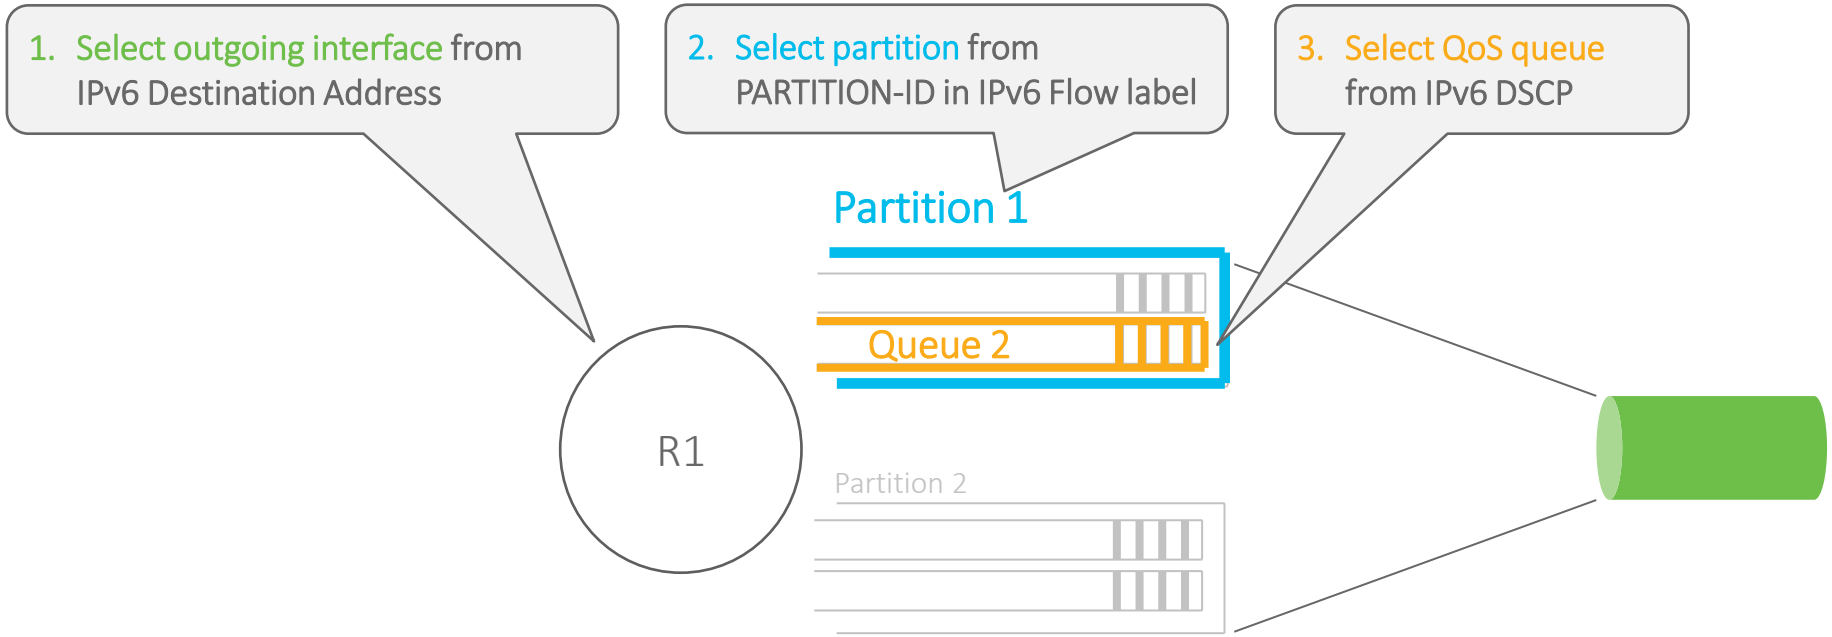

- Stateless: DA, PARTITION-ID and DSCP are independent fields in the packet header
- Scalable: Routing, Partitioning & QoS are orthogonal spaces
- Seamless Deployment: PARTITION-ID is seamlessly hashed by legacy devices
- IETF Terminology: Network Resource Partitioning (draft-filsfils-spring-srv6-stateless-slice-id)

Demo Link

SRv6 uSIDs

SRv6 uSID Terminology

- Industry:
  - SRv6 Micro Segment
  - SRv6 uSID
  - Briefly: uSID
- IETF: NEXT-C-SID
  - Briefly: Next
  - IETF document: draft-ietf-spring-srv6-srh-compression-01
  - Training: link

Rakuten – SRv6 uSID in Deployment

- 5G End-to-End Network Slicing based on SRv6 u-SID Flex-Algo
  - SRv6 uSIDs are allocated from the ULA address range
  - SRv6 uSID ISIS Flex-Algo: Low-Cost vs Low-Delay
  - SRv6 uSID BGP services
  - SRv6 uSID TILFA
  - Cisco NCS5500 and NCS-540 series
- Innovation in partnership with Cisco
  - BGP PIC Core and Edge with SRv6 Summarization: ISIS UPA (demo)
  - SR BW counters for deterministic and scalable capacity planning and BW guarantee

Rakuten end-to-end 5G Network Slicing using SRv6 uSID

Bell Canada - SRv6 uSID Deployment

- Bell promptly switched from SR-MPLS to SRv6 uSID
- Continued Simplification (remove MPLS dataplane)
- Better Routing Scale: Summarization
- Better HW Scale: linerate 26 uSID push for end-to-end SR Policy
- Seamless Deployment (6 uSID's in DA without SRH)
- End-to-End IP Unified Dataplane from socket to Internet Peering
  - SRv6-TE Policy: topological and service uSID's
- Service Programming
- Reduce network service costs by up to 90%
  footprint by 75%
  power consumption by as much as 66%

Bell SRv6 uSID Deployment

Daniel Bernier
Technical Director, Bell Canada

Jesper Eriksson,
VP Product Management, NoviFlow inc

SRv6 and P4 at the Network Edge

SRv6 uSID - Feature Parity with SR-MPLS

- TILFA & uLoop Avoidance
- Flex-Algo Low-Cost/Low-Delay with VPN Automated Steering
  - Performance Monitoring: Link Latency (for Low-Delay slice)
- SRv6-TE Policy: topological and service uSID’s
- L3VPN (IPv4 and IPv6), IPv4 Internet, IPv6 Internet, PW, BD
- VPN GW to interconnect with legacy VPN
- Seamless Inter-Domain with Summarization
  - No need for complex BGP3107 mechanisms

SRv6 uSID - Rich Eco-System

- Cisco, NoviFlow, Arrcus, Nokia
- Merchant: Silicon One, Broadcom, Marvell, Barefoot
- Open Source: Linux, FD.io, P4, eBPF, Cillium, SAI

SRv6 uSID – Pure IP Routing

- Classless Routing (CIDR, RFC7608)

“CIDR rules, even within an SR domain. For that reason, the fact that the bottom 64 bits in the "address" look funny or change is simply irrelevant. They are invisible to routing (which is done based on the prefix)....”

Brian Carpenter, former IETF and IAB chair

https://mailarchive.ietf.org/arch/msg/ipv6/37bH3Ag7jaNgKwnZY_mKIxWRwdM/

IP and uSID – Pure IP Forwarding

- Longest-Prefix-Match (LPM)
- Benefits
  - Straightforward Implementation
  - Support for different uSID sizes in the same uSID list
  - Capability to process multiple uSIDs with one single LPM
- Other technologies cannot do this
  - An MPLS lookup processes one single label at a time
  - All MPLS labels have a fixed 20-bit size
  - GSID (CSID-REPLACE) has the same limitations as MPLS

Perfect SRv6 Integration

- uSID reuses SRH (RFC8754) without any change
- uSID strictly applies the SRv6 Network Programming (RFC8986)
- uSID can be bound to any Network Programming Instruction

SRv6 uSID offers the best SRv6 Compression

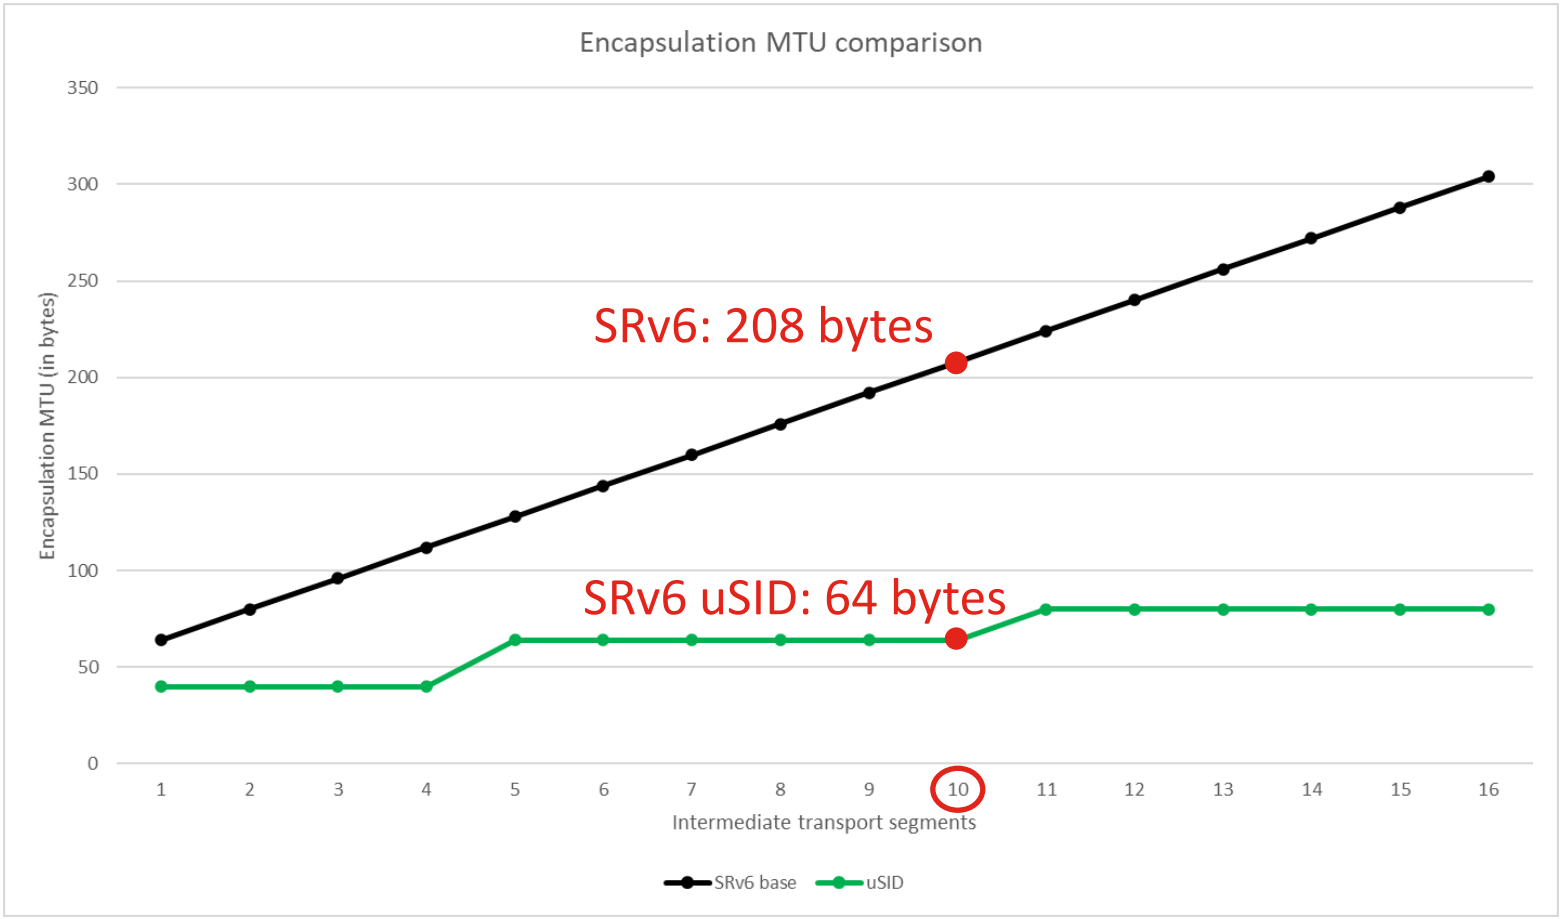

SRv6 uSID offers the best SRv6 Compression

uSID (NEXT):
average compression 62%

GSID (REPLACE):
average compression 39%

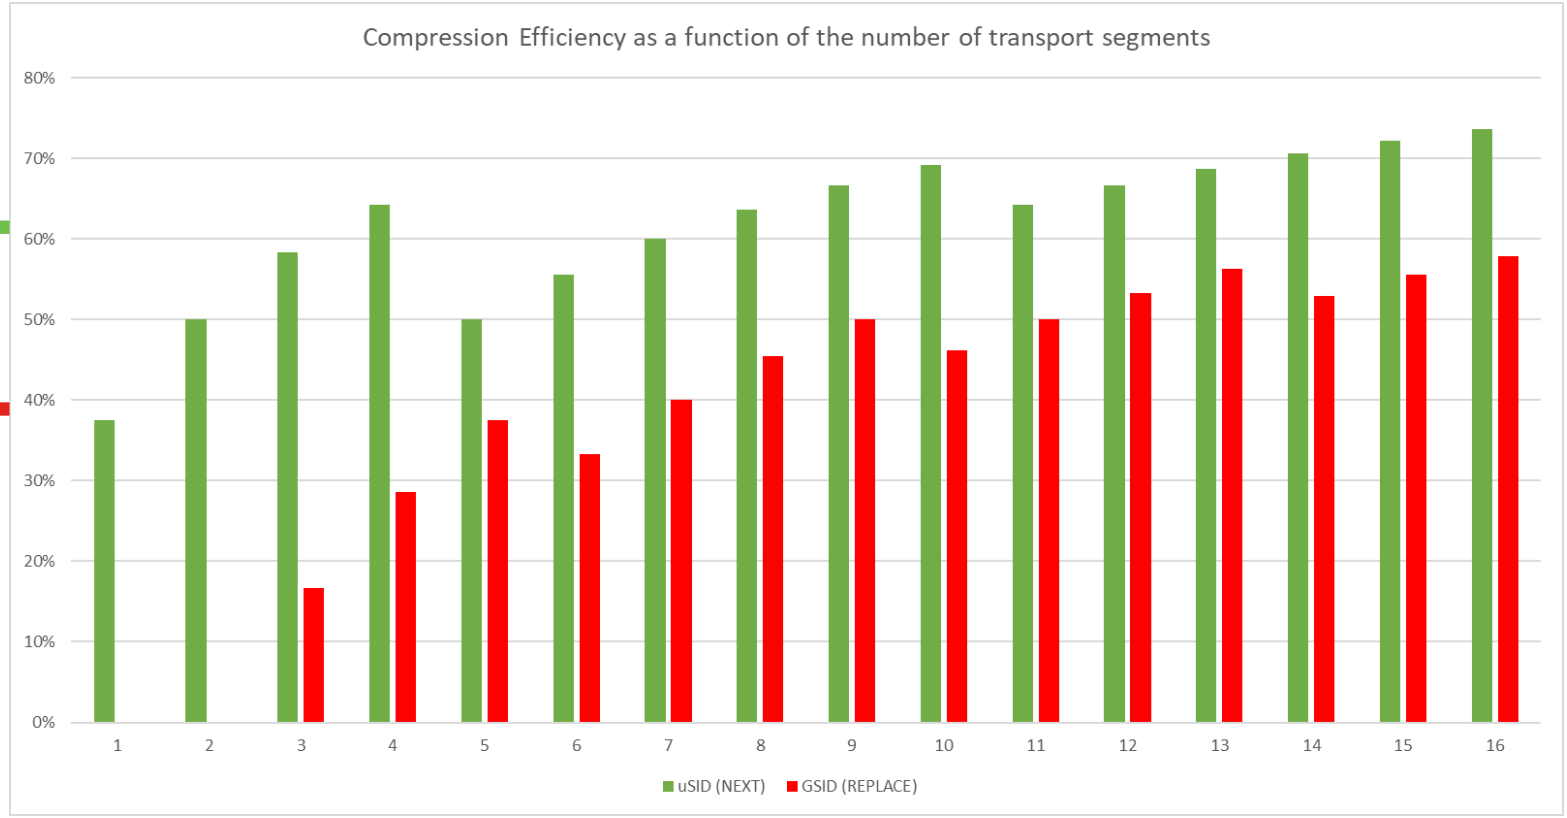

uSID is 59% more efficient than GSID (CSID-Replace)

62/39 = 1.59

Better HW Scale

	SRv6 uSID	MPLS
Linerate steering into SR Policy of N SID's (NCS-5700, J2)	N=26	N=~12
Consumed counters associated to a remote ISIS node	1	4
Consumed dataplane entries associated to remote ISIS node	1	4

Demo Link

Better Routing Scale

	SRv6 uSID	MPLS
Unique Nodes in the SR domain	15M-240M	0.8M
Unique Services per node	512k	200k
ISIS Summarization	Yes	No
BGP3107 complexity tax to scale ISIS Host Routes	No	Yes

Seamless Deployment

- Many use-cases leverage classic IPinIP
- The outer DA already holds 6 uSIDs !! D. Cai – Alibaba podcast - SRv6 uSID
- SDN coherence
  - The SDN controller packs the uSID policy in the outer DA
  - The host/access node only sees an opaque DA
- VPN GW to interconnect with legacy VPN

Path Tracing

The exact path from A to M is not known

- 7 possible “valid” ECMP path
  - ABFM, ABGM, ACFM, ACGM, ACHM, ADGM, ADHM …
- The path may be invalid
  - Routing or FIB corruption @ B
- Timestamp at each hop
- Interface Load at each hop

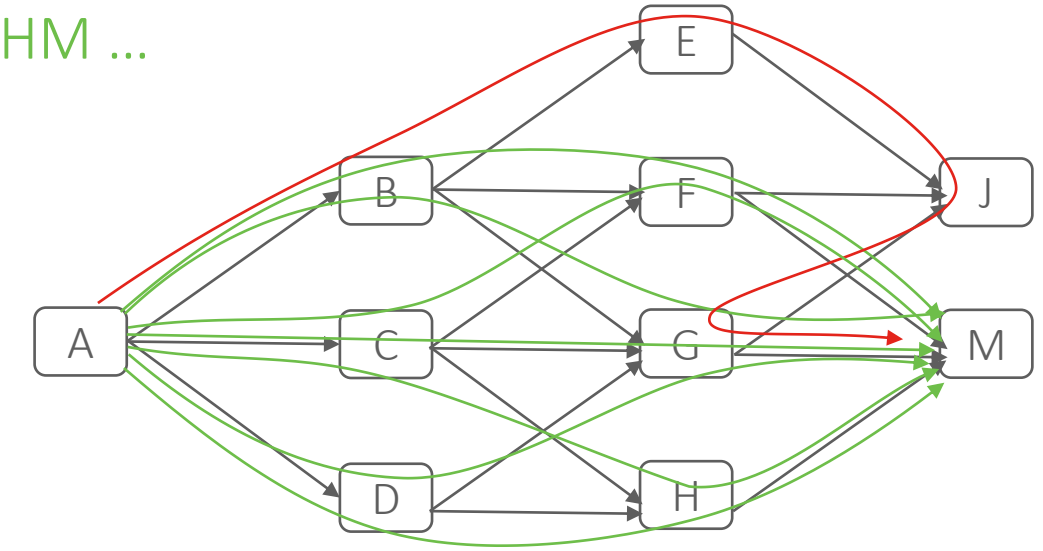

Stamping Trajectory in PT Header

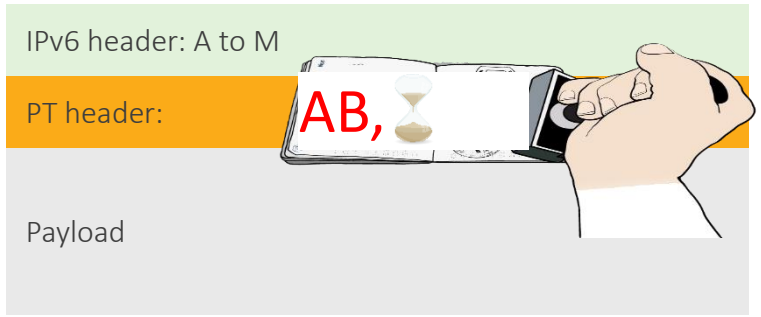

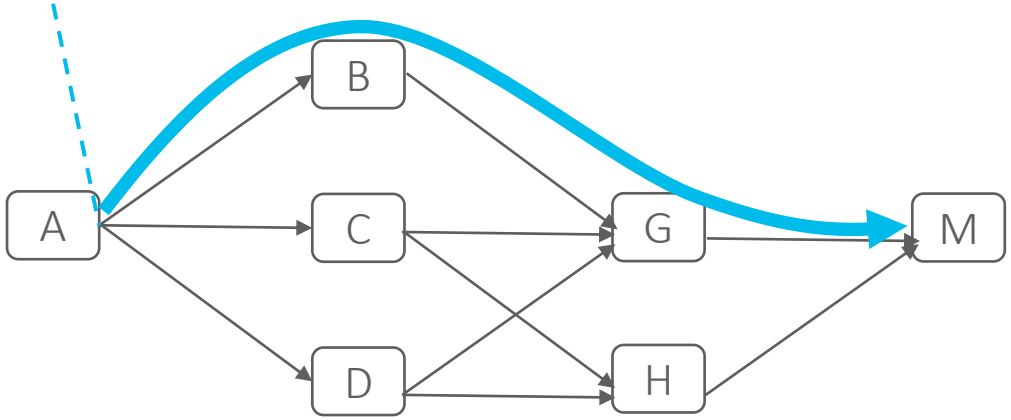

Stamping Trajectory in PT Header

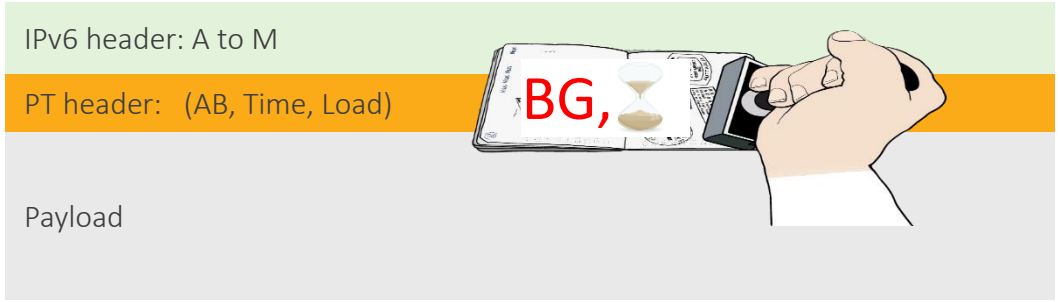

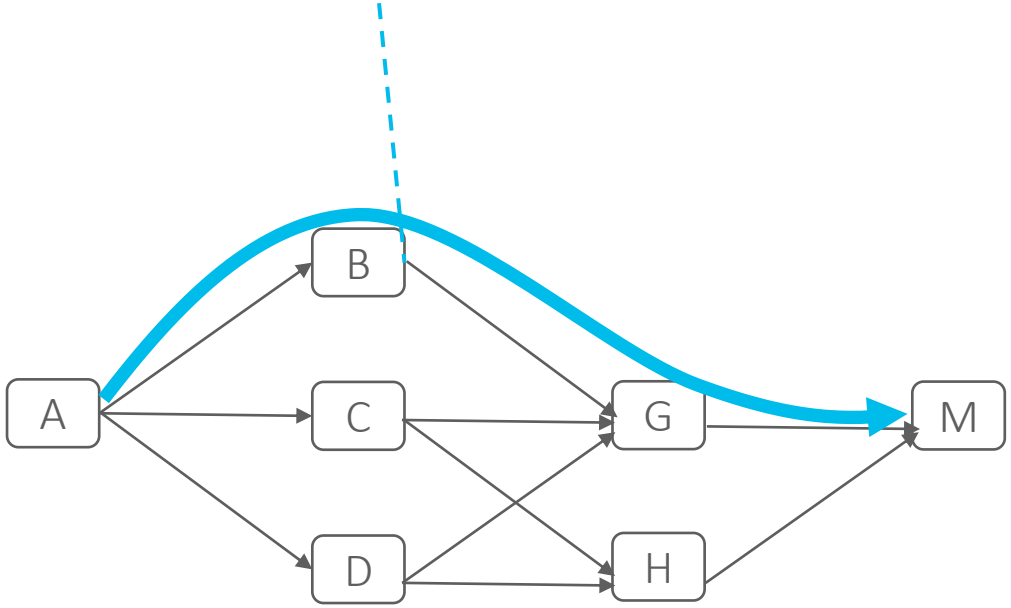

Stamping Trajectory in PT Header

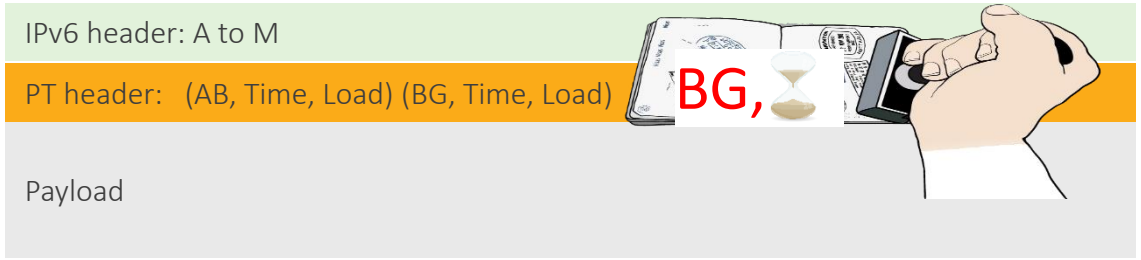

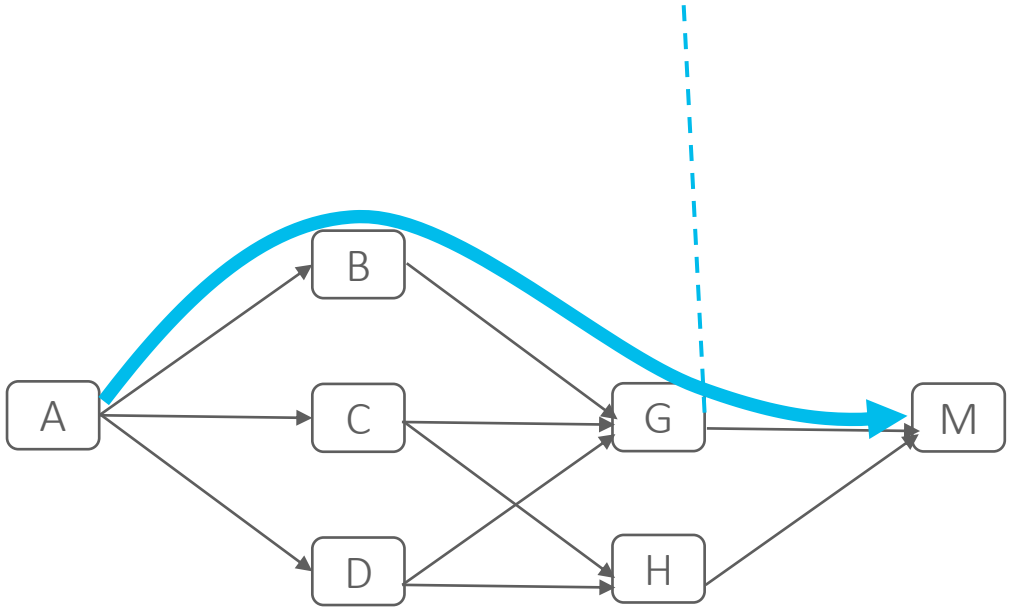

The PT idea

- Stamping in the Packet Header
- Implemented in the most basic HW pipeline
  - Linerate for any packet
- Ultra-MTU-efficient: only 3 bytes per hop!
  - 12-bit Interface, 8-bit Timestamp, 4-bit Load
- For IPv6, with or without SRH
  - MPLS solution also designed
- Interwork with legacy node

Product, Deployment & EcoSystem

- Cisco Shipping in CY22
  - PT Demo and Training (link)

- Strong Operator Interest
- Rich Eco-System

- Rich Open-Source

Conclusion

Simplicity Always Prevails

~~LDP~~

~~RSVP-TE~~

~~BGP 3108~~

~~MPLS~~

~~UDP/VxLAN~~

~~NSH~~

Furthermore with more scale and functionality

Stay up-to-date

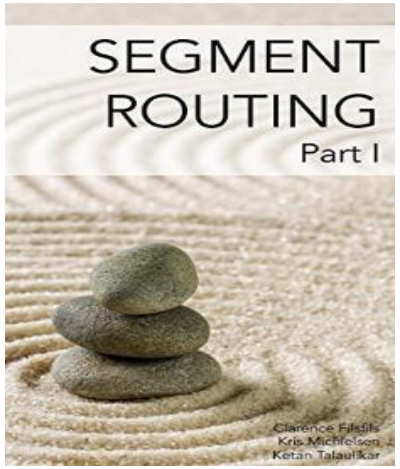

amzn.com/B01I58LSUO

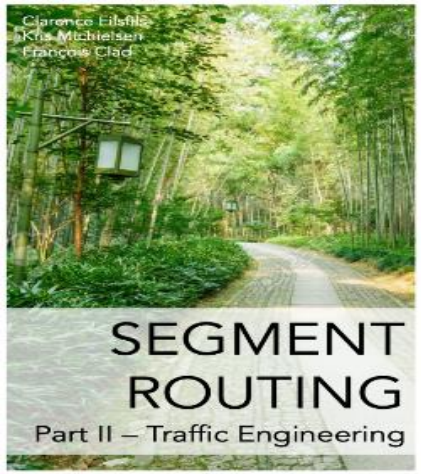

amazon.com/dp/B07N13RDM9

SRv6 Part III Coming by June 2022

twitter.com/SegmentRouting

SR Segment Routing

segment-routing.net

facebook.com/SegmentRouting/

linkedin.com/groups/8266623

ask-segment-routing@cisco.com

Simplicity Always Prevails

~~LDP~~

~~RSVP-TE~~

~~BGP 3108~~

~~MPLS~~

~~UDP/VxLAN~~

~~NSH~~

Furthermore with more scale and functionality

Appendix

SRv6 – Engineered for Simplicity at ANY Scale

- Unified IP Dataplane: from socket to Internet peering though DC, Access, Metro, Core…
- Stateless
  - SID list is in the header, Slice ID is in the header
- Automated
  - TILFA, uLoop, BGP PIC, Automated Steering, ODN SR Policy
- Simplification
  - No MPLS Control-Plane: LDP, RSVP-TE and MPLS OAM
  - No MPLS Data-Plane: FIB entries and counters (ip2mpls, mpls2ip, mpls2mpls), ucode
  - Fast Routing Convergence: less entries to update
  - No need for BGP3107 to distribute MPLS host-routes
- Optimum Load-Balancing
  - MPLS never solved this correctly
- Better Routing Scale
  - Far more Global and Local ID’s
  - ISIS/OSPF Summarization
- Seamless Deployment: SRH rarely needed (6 uSID’s in DA), transit routers are pure IPv6
- Optimal MTU Efficiency: SRv6 uSID
- Enhanced Data Plane Monitoring (Path Tracing)
- Performance monitoring (Latency, Availability, Loss)

Same Security Model as any VPN solution

- Border Routers block any traffic destined to the VPN ID’s within the domain
  - VxLAN
  - L2TPv3
  - GRE
  - SRv6
- Additional protection
  - Allocate the SID’s from unroutable FC/7 block

Impressive 5G Deployments

News Release

Rakuten Mobile Advances Its Network for 5G and IoT Services with Cisco

June 25, 2021

TOKYO, JAPAN, June 25, 2021 – Cisco and Rakuten Mobile, Inc. today announced a major milestone for Rakuten Mobile's network infrastructure in support of efforts to build a better, more inclusive internet for the future.

Rakuten Mobile operates the world first's fully cloud-native mobile network. It launched 4G service in Japan in April 2020, and launched 5G non-standalone (NSA) services in September 2020 in record time. With four million subscribers today, Rakuten Mobile continues to advance and scale its network to support new demands driven by the growth of remote and mobile workers.

With the implementation of Segment Routing over IPv6 (SRv6) and Cisco Routed Optical Networking, Rakuten Mobile plans to expand its capabilities to support enterprise customers with 5G and IoT services. To support its future 5G SA services with network slicing capabilities, Rakuten Mobile will introduce SRv6 micro-segments, an extension to the SRv6 network programming model that is key to addressing multi-domain 5G deployments

Impressive SDN Deployments

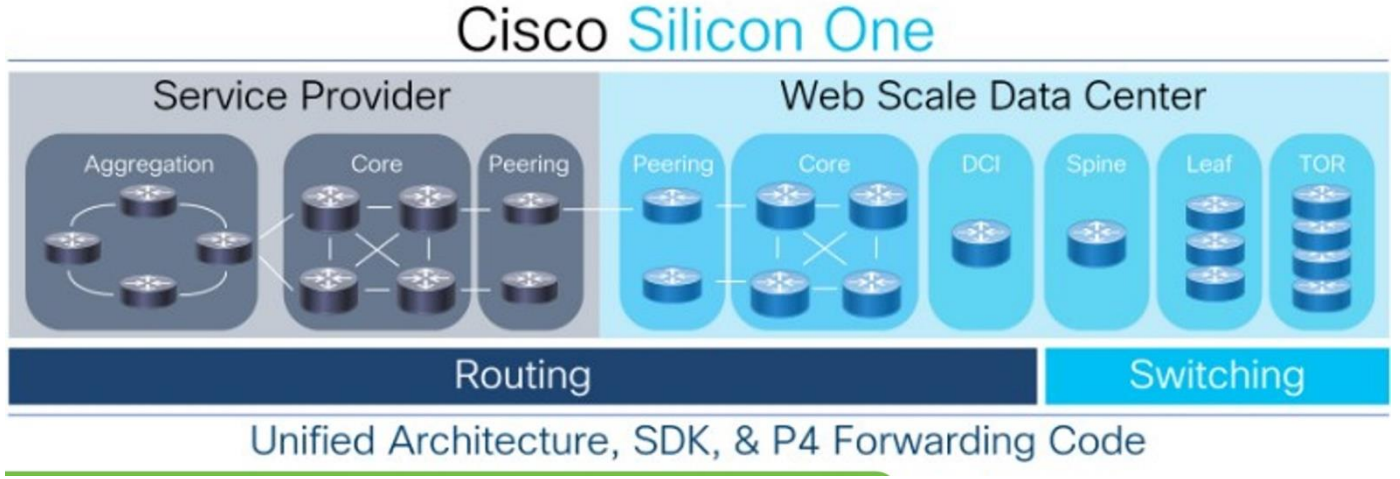

"Embracing the path towards network softwarization and "in-network" or ubiquitous computing, SRv6 Network Programming and its uSID implementation brings software like programmability and agility at the protocol level, with a broad set of implementations.

Silicon1 and its P4 SDK brings it a step further by extending its reach at the ASIC level for massive scale, speed and cost efficiency. "

Daniel Bernier


CISCO